What is claimed is:

1. A method of determining a placement of services of a distributed application onto nodes of a distributed resource infrastructure comprising the steps of:
   a. establishing a placement indicator for a specific service;
   b. forming communication constraints between node pairs which ensure that a sum of transport demands between a particular node pair does not exceed a transport capacity between the particular node pair, each term of the sum comprising a product of a first placement variable, a second placement variable, and the transport demand between the services associated with the first and second placement variables;
   c. forming an objective; and
   d. employing a local search solution to solve an integer program comprising the placement indicator, the communication constraints, and the objective to determine the placement of the services onto the nodes.

2. The method of claim 1 wherein the placement indicator comprises a pre-defined placement.

3. The method of claim 2 wherein the pre-defined placement comprises placing the specific service onto a specific node.

4. The method of claim 2 wherein the pre-defined placement comprises not placing the specific service onto a specific node.

5. The method of claim 1 wherein the placement indicator comprises a neutral indication of whether the specific service is to be placed onto a specific node.

6. A method of determining a placement of services of a distributed application onto nodes of a distributed resource infrastructure comprising the steps of:
   a. establishing an application model of the services comprising transport demands between the services;
   b. establishing an infrastructure model of the nodes comprising transport capacities between the nodes;

c. establishing a placement model comprising placement indicators for the services;

d. forming an integer program that comprises:

i. a set of placement variables for a combination of the services and the nodes, each of the placement variables indicating whether a particular service is located on a particular node;

ii. communication constraints between node pairs which ensure that a sum of the transport demands between a particular node pair does not exceed the transport capacity between the particular node pair, each term of the sum comprising a product of a first placement variable, a second placement variable, and the transport demand between the services associated with the first and second placement variables;

iii. placement constraints for the services which ensure that the services are placed onto the nodes in accord with the placement indicators; and

iv. an objective; and

e. employing a local search solution to solve the integer program which determines the placement of the services onto the nodes.

7. The method of claim 6 wherein a particular placement indicator comprises an indication that a specific service is to be placed onto a specific node.

8. The method of claim 6 wherein a particular placement indicator comprises an indication that a specific service is not to be placed onto a specific node.

9. The method of claim 6 wherein a particular placement indicator comprises a neutral indication of whether a specific service is to be placed onto a specific node.

10. The method of claim 9 wherein a default for the placement indicators comprises the neutral indication.

11. A method of determining a placement of services of a distributed application onto nodes of a distributed resource infrastructure comprising the steps of:

a. establishing an application model of the services that comprises processing

demands for the services, storage demands for the services, and transport demands between the services;

b. establishing an infrastructure model of the nodes that comprises processing capacities for the nodes, storage capacities for the nodes, and transport capacities between the nodes;

c. establishing a placement model comprising placement indicators for the services;

d. forming an integer program that comprises:

i. a set of placement variables for a combination of the services and the nodes, each of the placement variables indicating whether a particular service is located on a particular node;

ii. processing constraints which ensure that a sum of the processing demands for each of the nodes does not exceed the processing capacity for the node;

iii. storage constraints which ensure that a sum of the storage demands for each of the nodes does not exceed the storage capacity for the node;

iv. first placement constraints which ensure that each of the services is placed on one and only one node;

v. second placement constraints which ensure that the services are placed onto the nodes in accord with the placement indicators;

vi. communication constraints between node pairs which ensure that a sum of the transport demands between a particular node pair does not exceed the transport capacity between the particular node pair, each term of the sum comprising a product of a first placement variable, a second placement variable, and the transport demand between the services associated with the first and second placement variables; and

vii. an objective of minimizing communication traffic between the nodes and balancing processing loads on the nodes; and

e. employing a local search solution to solve the integer program which determines the placement of the services onto the nodes.

12. A computer readable memory comprising computer code for directing a computer to make a determination of a placement of services of a distributed application onto nodes of a distributed resource infrastructure, the determination

of the placement of the services onto the nodes comprising the steps of:

a. establishing a placement indicator for a specific service;

b. forming communication constraints between node pairs which ensure that a sum of transport demands between a particular node pair does not exceed a transport capacity between the particular node pair, each term of the sum comprising a product of a first placement variable, a second placement variable, and the transport demand between the services associated with the first and second placement variables;

c. forming an objective; and

d. employing a local search solution to solve an integer program comprising the placement indicator, the communication constraints, and the objective to determine the placement of the services onto the nodes.

13. The computer readable memory of claim 12 wherein the placement indicator comprises a pre-defined placement.

14. The computer readable memory of claim 13 wherein the pre-defined placement comprises placing the specific service onto a specific node.

15. The computer readable memory of claim 13 wherein the pre-defined placement comprises not placing the specific service onto a specific node.

16. The computer readable memory of claim 12 wherein the placement indicator comprises a neutral indication of whether the specific service is to be placed onto a specific node.

17. A computer readable memory comprising computer code for directing a computer to make a determination of a placement of services of a distributed application onto nodes of a distributed resource infrastructure, the determination of the placement of the services onto the nodes comprising the steps of:

a. establishing an application model of the services comprising transport demands between the services;

b. establishing an infrastructure model of the nodes comprising transport capacities between the nodes;

c. establishing a placement model comprising placement indicators for the services;

d. forming an integer program that comprises:

i. a set of placement variables for a combination of the services and the nodes, each of the placement variables indicating whether a particular service is located on a particular node;

ii. communication constraints between node pairs which ensure that a sum of the transport demands between a particular node pair does not exceed the transport capacity between the particular node pair, each term of the sum comprising a product of a first placement variable, a second placement variable, and the transport demand between the services associated with the first and second placement variables;

iii. placement constraints for the services which ensure that the services are placed onto the nodes in accord with the placement indicators; and

iv. an objective; and

e. employing a local search solution to solve the integer program which determines the placement of the services onto the nodes.

18. The computer readable memory of claim 17 wherein a particular placement indicator comprises an indication that a specific service is to be placed onto a specific node.

19. The computer readable memory of claim 17 wherein a particular placement indicator comprises an indication that a specific service is not to be placed onto a specific node.

20. The computer readable memory of claim 17 wherein a particular placement indicator comprises a neutral indication of whether a specific service is to be placed onto a specific node.

21. The computer readable memory of claim 20 wherein a default for the placement indicators comprises the neutral indication.

22. The computer readable memory of claim 20 wherein a matrix is specified

which expresses constraints or preferences for identifying a placement of services onto nodes.

23. A computer readable memory comprising computer code for directing a computer to make a determination of a placement of services of a distributed application onto nodes of a distributed resource infrastructure, the determination of the placement of the services onto the nodes comprising the steps of:

a. establishing an application model of the services that comprises processing demands for the services, storage demands for the services, and transport demands between the services;

b. establishing an infrastructure model of the nodes that comprises processing capacities for the nodes, storage capacities for the nodes, and transport capacities between the nodes;

c. establishing a placement model comprising placement indicators for the services;

d. forming an integer program that comprises:

i. a set of placement variables for a combination of the services and the nodes, each of the placement variables indicating whether a particular service is located on a particular node;

ii. processing constraints which ensure that a sum of the processing demands for each of the nodes does not exceed the processing capacity for the node;

iii. storage constraints which ensure that a sum of the storage demands for each of the nodes does not exceed the storage capacity for the node;

iv. first placement constraints which ensure that each of the services is placed on one and only one node;

v. second placement constraints which ensure that the services are placed onto the nodes in accord with the placement indicators;

vi. communication constraints between node pairs which ensure that a sum of the transport demands between a particular node pair does not exceed the transport capacity between the particular node pair, each term of the sum comprising a product of a first placement variable, a second placement variable, and the transport demand between the services associated with the first and second placement variables; and

vii. an objective of minimizing communication traffic between the nodes and balancing processing loads on the nodes; and

e. employing a local search solution to solve the integer program which determines the placement of the services onto the nodes.